Differential Game Logic for Hybrid Games

André Platzer

March 2012
CMU-CS-12-105

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

This material is based upon work supported by the National Science Foundation under NSF CAREER Award CNS-1054246 and NSF EXPEDITION CNS-0926181 and by the Army Research Office under Award No. W911NF-09-1-0273. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution or government. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of any ponsoring institution or government.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
MAR 2012		**00-00-2012 to 00-00-2012**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Differential Game Logic for Hybrid Games	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Carnegie Mellon University,School of Computer Science,Pittsburgh,PA,15213**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

We introduce differential game logic (dGL) for specifying and verifying properties of hybrid games, i.e., determined, sequential/dynamic, non-cooperative, zero-sum games of perfect information on hybrid systems that combine discrete and continuous dynamics. Unlike hybrid systems hybrid games allow choices in the system dynamics to be resolved by different players with different objectives. The logic dGL can be used to study properties of the resulting adversarial behavior. It unifies differential dynamic logic for hybrid systems with game logic. We define a regular modal semantics for dGL, present a proof calculus for dGL, and prove soundness. We identify separating axioms, i.e., the axioms that distinguish dGL and its game aspects from logics for hybrid systems. We also define an operational game semantics, prove equivalence, and prove determinacy.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**Same as Report (SAR)**	**27**	

Keywords: dynamic logic, game logic, hybrid games, hybrid dynamical systems, proof calculus

Abstract

We introduce *differential game logic* (dGL) for specifying and verifying properties of *hybrid games*, i.e., determined, sequential/dynamic, non-cooperative, zero-sum games of perfect information on hybrid systems that combine discrete and continuous dynamics. Unlike hybrid systems, hybrid games allow choices in the system dynamics to be resolved by different players with different objectives. The logic dGL can be used to study properties of the resulting adversarial behavior. It unifies differential dynamic logic for hybrid systems with game logic. We define a regular modal semantics for dGL, present a proof calculus for dGL, and prove soundness. We identify separating axioms, i.e., the axioms that distinguish dGL and its game aspects from logics for hybrid systems. We also define an operational game semantics, prove equivalence, and prove determinacy.

1 Introduction

Hybrid systems [Hen96] are dynamical systems that combine discrete dynamics and continuous dynamics. They are important for modeling systems that use computers to control physical systems. Hybrid systems allow discrete jump assignments for discrete dynamics and differential equations for continuous dynamics. They combine conditional switching, nondeterminism, and repetition. Hybrid systems are undecidable [Hen96, AM98, CL00], but nevertheless the focus of many successful verification approaches. They have a complete axiomatization relative to differential equations in a logic called *differential dynamic logic* (dL) [Pla08, Pla12b, Pla12a], which extends Pratt's dynamic logic of conventional discrete programs [Pra76] to hybrid systems.

In this paper, we consider multi-agent hybrid systems, where two agents act and we are uncertain how they will interact with each other. Agents often have only limited knowledge about their environment or about the exact future behavior of other agents. In that case, the system turns into a game in which every agent has a set of actions to choose from as the system evolves. Each agent can control its own actions to realize its own objective but has to be prepared to handle *all* possible actions by other agents who may follow other objectives. Because the agents play on a hybrid system, we obtain a *hybrid game*, i.e., a game of two agents on a hybrid system [TPS98, TLS00, VPVD11]. Hybrid systems also allow for nondeterminism and previous logics can be used to prove properties about all (dL formula $[\alpha]\phi$) or some ($\langle\alpha\rangle\phi$) ways of resolving it [Pla08]. In hybrid systems, exactly one entity chooses how to resolve the nondeterminism. In hybrid games, instead, two players have the opportunity to resolve nondeterministic choices interactively, based on the outcome that previous decisions by the other player have had. Hybrid games are sequential/dynamic, noncooperative zero-sum two-player games of perfect information played on hybrid systems. They are based on discrete games [vNM55, Nas51], which have been studied more exhaustively. Zero-sum two-player games are general in that any non-zero sum n-player game reduces to a zero-sum $(n+1)$-player game [vNM55, 56.2.2], and any n-player zero-sum game can be based on zero-sum two-player games of a player against an aggregate player [vNM55, 25.2]. Note that, even if the agents do not necessarily actively pursue the interest to spoil each others' objectives, they may still do so out of ignorance, or because their respective actions interfere. Every agent, thus, has to choose his actions in *some* way while being prepared that other agents could choose *any* of their actions, which is an adversarial resolution of the nondeterminisms in the game.

Games and logic have been shown to interact fruitfully in many ways [HS97]. We focus on using logic to specify and verify properties of hybrid games. Our approach to verifying hybrid games is inspired by Parikh's game logic [Par85, PP03]. Game logic generalizes (propositional discrete) dynamic logic to discrete games played on a finite state spaces. We introduce a logic, *differential game logic* (dGL), that generalizes differential dynamic logic (dL) [Pla08, Pla12b, Pla12a] to hybrid games and, simultaneously, generalizes game logic [Par85, PP03] to hybrid systems with their uncountable state spaces and interacting discrete and continuous dynamics.

The logic dGL we present here has some similarity with our *stochastic differential dynamic logic* (SdL) [Pla11], because both address the issue of how to verify properties of the system dynamics with partially uncertain behavior. Both approaches do, however, address uncertainty in fundamentally different ways. SdL takes a probabilistic perspective on uncertainty in the system dynamics. The dGL approach put forth in this paper, instead, takes an adversarial perspective

on uncertainty. Both views on how to handle uncertain behavior are useful but serve different purposes, depending on the nature of the system analysis question at hand. A probabilistic understanding of uncertainty can be superior whenever good information is available about the distribution of choices made by the environment. Whenever that is not possible, adversarial views are more appropriate, since they do not lead to the inadequate biases that arbitrary probabilistic assumptions would impose. Security questions about hybrid systems lead to inherently adversarial situations. Controller synthesis for hybrid systems is another application that reduces to a hybrid game [VPVD11].

Our primary contributions are that we identify the logical essentials of hybrid games and their game combinators, introduce differential game logic, a semantics, and proof calculus, and that we characterize what constitutes the fundamental difference of hybrid systems proving compared to hybrid games proving. Furthermore, we relate this semantics to a game-theoretical operational game semantics, prove equivalence, and prove determinacy.

2 Differential Game Logic

The games we consider have no draws and if a player is deadlocked, he loses. If the game completes without deadlock, the player who reaches one of his winning states wins. Thus, exactly one player wins each game play, since the winning states are complementary. Our games are zero-sum games, i.e., if one player wins, the other one loses, and vice versa, with player payoffs ± 1. Classically, the two players are called *Angel* and *Demon*. Our games are non-cooperative and sequential games. That is, the players do not negotiate binding contracts (beyond what is represented in the rules of the game), but can choose to act at will. Furthermore, the games are sequential (or dynamic), i.e., the game proceeds in a series of steps. At each step, exactly one of the players can choose an action and his next action can be based on the outcome of the last action (by the other player or himself, whoever moved last) and, thus, may depend on the previous choices determining the current state.

The *hybrid games of differential game logic* dGL are defined by the following grammar (α, β are hybrid games, x a vector of variables, θ a vector of terms of the same dimension, H a formula of first-order arithmetic, and ϕ is a dGL formula, usually first-order):

$$\alpha, \beta \ ::= \ x := \theta \mid ?\phi \mid x' = \theta \,\&\, H \mid \alpha \cup \beta \mid \alpha; \beta \mid \alpha^* \mid \alpha^d$$

The *formulas of differential game logic* dGL are defined by the following grammar (ϕ, ψ are dGL formulas, θ_i are terms, x a variable, and α is a hybrid game):

$$\phi, \psi \ ::= \ \theta_1 \geq \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \exists x\, \phi \mid \langle\alpha\rangle\phi$$

The operator $[\alpha]$ dual to $\langle\alpha\rangle$ is defined by $[\alpha]\phi \equiv \neg\langle\alpha\rangle\neg\phi$. Operators $>, =, \leq, <, \vee, \rightarrow, \leftrightarrow, \exists x$ can be defined as usual, e.g., $\forall x\, \phi \equiv \neg\exists x\, \neg\phi$. Formula $\langle\alpha\rangle\phi$ expresses that Angel has a winning strategy to achieve ϕ in game α, i.e., Angel has a strategy to reach a state satisfying formula ϕ when playing game α, no matter what strategy Demon chooses. The formula $[\alpha]\phi$ expresses that Angel does not have a winning strategy to achieve $\neg\phi$ in game α. This is equivalent to Demon having a

winning strategy to achieve ϕ, because, any way how Demon plays to prevent Angel from winning is a winning strategy for Demon, since there are no draws and the game cannot be played infinitely long. That is, our games are *determined*, i.e., from each state and for each winning condition ϕ, either Angel has a winning strategy or Demon has a winning strategy. Determinacy follows from the Borel determinacy theorem [Kec94, Theorem 20.6]; see Section 4 for details.

The atomic games of dGL are assignments, continuous evolutions, and tests. In the *deterministic assignment game* $x := \theta$, the value of variable x changes instantly and deterministically to that of θ without any choice to resolve. In the *continuous evolution game* $x' = \theta \,\&\, H$, the duration of the evolution of the continuous evolution along differential equation $x' = \theta$ is Angel's choice, but Angel is not allowed to choose a duration that would cause the state to leave the region where formula H holds. In particular, Angel is deadlocked and loses if H does not hold in the current state, because she cannot even evolve for duration 0 then. The *test game* or *challenge* $?\phi$ has no effect on the state, except that Angel loses the game if dGL formula ϕ does not hold in the current state.

The compound games are sequential composition, choice, repetition, and duals. The *sequential game* $\alpha;\beta$ is the game that first plays game α and, when game α terminates without a player having won already, continues by playing game β. In the *choice game* $\alpha \cup \beta$, Angel chooses whether to play game α or play game β. The *repeated game* α^* plays game α repeatedly and Angel chooses, after each play of α that terminates without a player having won already, whether to play the game again or not, but she cannot choose to play infinitely often (any number $n \in \mathbb{N}$ of repetitions is permitted, including zero). Thus, we consider games on non-Zeno hybrid system runs [DN00, Hen96]. The *dual game* α^d is the same as playing the game α with the roles of the players swapped. That is, in α^d, Demon decides all choices that Angel has in α, and Angel decides all choices in α^d that Demon has in α. Players who are supposed to move but deadlock lose. Test game $?\phi$ causes Angel to lose if formula ϕ does not hold. Dual test game $(?\phi)^d$ causes Demon to lose if ϕ does not hold.

Demonic choice between game α and β is $(\alpha^d \cup \beta^d)^d$ and denoted by $\alpha \cap \beta$, in which either the game α or the game β is played, by Demon's choice. *Demonic repetition* of game α is $((\alpha^d)^*)^d$ and denoted by $\alpha^\times$, in which α is repeated as often as Demon chooses to. In $\alpha^\times$, Demon chooses after each play of α whether to repeat the game, but cannot play infinitely often. The dual operator ^{d} is the only syntactic difference of dGL for hybrid games compared to dL for hybrid systems [Pla08, Pla12b, Pla12a], but a fundamental one, because it is the only operator where control passes from Angel to Demon or back. The *dual differential equation* $(x' = \theta \,\&\, H)^d$ follows the same dynamics as $x' = \theta \,\&\, H$ except that Demon chooses the duration. Dual assignment $(x := \theta)^d$ is equivalent to $x := \theta$, because it involves no choices.

Observe that every play of a game is won or lost by exactly one player. Even a repeated game α^* has only one winner, because the game stops as soon as one player has won. This is different than the classical repetition of game plays (including winning/losing), where the purpose is for the players to repeat the same game over and over again, win and lose multiple times, and study who wins how often in the long run with mixed strategies. In our scenario, the overall game is played once (even if some part of it constitutes in repeating action choices) and stops as soon as either Angel or Demon have won. In applications, the system is already in trouble even if it loses the

game only once, because that may entail that a safety-critical property has already been violated.

3 Semantics

A *state* s is a mapping from variables to $\mathbb{R}$. The set of states is denoted by $\mathcal{S}$ and isomorphic to a Euclidean space $\mathbb{R}^n$ when n is the number of variables. We use s_x^d to denote the state that agrees with state s except for the interpretation of variable x, which is changed to $d \in \mathbb{R}$. We denote the value of term θ in s by $[\![\theta]\!]_s$. The *semantics of a* dGL *formula* ϕ is the subset $[\![\phi]\!] \subseteq \mathcal{S}$ of states in which ϕ is true. It is defined as follows

1. $[\![\theta_1 \geq \theta_2]\!] = \{s \in \mathcal{S} : [\![\theta_1]\!]_s \geq [\![\theta_2]\!]_s\}$
2. $[\![\neg\phi]\!] = \mathcal{S} \setminus [\![\phi]\!]$
3. $[\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!]$
4. $[\![\exists x\, \phi]\!] = \{s \in \mathcal{S} : s_x^r \in [\![\phi]\!] \text{ for some } r \in \mathbb{R}\}$
5. $[\![\langle\alpha\rangle\phi]\!] = \varsigma_\alpha([\![\phi]\!])$

A dGL formula ϕ is *valid*, written $\vDash \phi$, iff $[\![\phi]\!] = \mathcal{S}$. The semantics of a hybrid game is not a reachability relation of a hybrid system, because the interactions of the players have to be taken into account. The *semantics of a hybrid game* α is a function $\varsigma_\alpha(\cdot)$ that, for each set of Angel's winning states $X \subseteq \mathcal{S}$ gives the set of states $\varsigma_\alpha(X)$ from which Angel has a winning strategy to achieve X (whatever strategy Demon chooses). It is defined as follows

1. $\varsigma_{x:=\theta}(X) = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in X\}$
2. $\varsigma_{x'=\theta\,\&\,H}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X \text{ for some } 0 \leq r \in \mathbb{R} \text{ and some (differentiable) } \varphi : [0,r] \to \mathcal{S} \text{ such that } \frac{\mathsf{d}\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)} \text{ and } \varphi(\zeta) \in [\![H]\!] \text{ for all } 0 \leq \zeta \leq r\}$
3. $\varsigma_{?\phi}(X) = [\![\phi]\!] \cap X$
4. $\varsigma_{\alpha\cup\beta}(X) = \varsigma_\alpha(X) \cup \varsigma_\beta(X)$
5. $\varsigma_{\alpha;\beta}(X) = \varsigma_\alpha(\varsigma_\beta(X))$
6. $\varsigma_{\alpha^*}(X) = \bigcap\{Z \subseteq \mathcal{S} : X \cup \varsigma_\alpha(Z) \subseteq Z\}$
7. $\varsigma_{\alpha^d}(X) = \mathcal{S} \setminus \varsigma_\alpha(\mathcal{S} \setminus X)$

Strategies do not occur explicitly in the dGL semantics, because it is based on the existence of winning strategies, not the strategies themselves. The semantics is fully compositional, i.e., the semantics of a compound dGL formula is a simple function of the semantics of its pieces, and the semantics of a compound hybrid game is a function of the semantics of its pieces. In particular, existence of a strategy in game α to achieve X is independent of any game and dGL formula

surrounding α, but just depends on the remaining game α itself and on the goal X. By an inductive argument, this reproves the classical result that we can focus on memoryless strategies, because the existence of strategies does not depend on surroundings, hence, by working bottom up, the strategy itself cannot depend on past states and choices, only the current state, remaining game, and goal.

Monotonicity, i.e., $\varsigma_\alpha(X) \subseteq \varsigma_\alpha(Y)$ for all $X \subseteq Y$, is easy to check for each case. Hence, the least fixpoint in $\varsigma_{\alpha^*}(X)$ is well-defined. The equivalence $[\alpha]\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi$ has two interesting consequences. The direction $\neg(\langle\alpha\rangle\phi \wedge [\alpha]\neg\phi)$ expresses that the game is *consistent*, i.e., from any state, at most one of the players can have a winning strategy for complementary winning conditions ϕ and $\neg\phi$, respectively. The direction $\langle\alpha\rangle\phi \vee [\alpha]\neg\phi$ represents that the game is *determined*, i.e., from any state, at least one of the players has a winning strategy to achieve complementary winning conditions ϕ and $\neg\phi$, respectively; see Section 4.

Note that dGL games branch finitely when the players decide which game to play in $\alpha \cup \beta$ and $\alpha \cap \beta$, respectively. The games α^* and $\alpha^\times$ also branch finitely, because, after each repetition of α, the respective player (Angel for α^* and Demon for $\alpha^\times$) may decide whether to repeat again or stop. Repeated games still lead to countably infinitely many branches, because a repeated game can be repeated any natural number of times. The game branches uncountably infinitely, however, when the players decide how long to evolve along differential equations in $x' = \theta \,\&\, H$ and $(x' = \theta \,\&\, H)^d$, because uncountably many nonnegative real number could be chosen as a duration (unless the system leaves H immediately).

In $\langle\alpha^*\rangle\phi$, Demon already has a winning strategy if he only has a strategy that prevents ϕ indefinitely, because Angel eventually has to stop repeating. Dually, in $\langle\alpha^\times\rangle\phi \equiv [\alpha^*]\phi$, Angel already has a winning strategy if she has a strategy that prevents ϕ indefinitely, because Demon eventually has to stop repeating.

Note that it is crucial that we have chosen finite repetition by the *least* fixpoint for the semantics of α^*. Otherwise, the *filibuster formula* would not have a well-defined truth-value:

$$\langle(x := 0 \cap x := 1)^*\rangle x = 0$$

The game in this formula never deadlocks (stalemates), because every player always has a remaining move (here even two). But, without the least fixpoint, the game would have perpetual checks, because no strategy helps either player win the game; see Fig. 1. Demon can move $x := 1$ and would win, but Angel observes this and decides to repeat, upon which Demon can again move $x := 1$. Thus (unless Angel is lucky starting from an initial state where she has won already) every strategy that one player has to reach $x = 0$ or $x = 1$ could be spoiled by the other player and the game would not be determined. Every player can let his opponent win, but would not have a strategy to win himself. Because of the least fixpoint $\varsigma_{\alpha^*}(X) = \mu Z.X \cup \varsigma_\alpha(Z)$ in the semantics, however, repetitions have to stop eventually (after an arbitrary and unbounded but finite number of rounds). That is why, in the example in Fig. 1, Demon wins and the formula is *false*, unless $x = 0$ already holds initially. Likewise, the dual filibuster game formula $x = 0 \to \langle(x := 0 \cup x := 1)^\times\rangle x = 0$ is (determined and) valid in dGL, because Demon has to stop repeating eventually.

Lemma 1 (Scott-continuity of non-interactive dGL)**.** *For ^{d}-free α, the semantics is Scott-continuous, i.e., $\varsigma_\alpha(\bigcup_{n\in I} X_n) = \bigcup_{n\in I} \varsigma_\alpha(X_n)$ for all families $\{X_n\}_{n\in I}$ with index set I.*

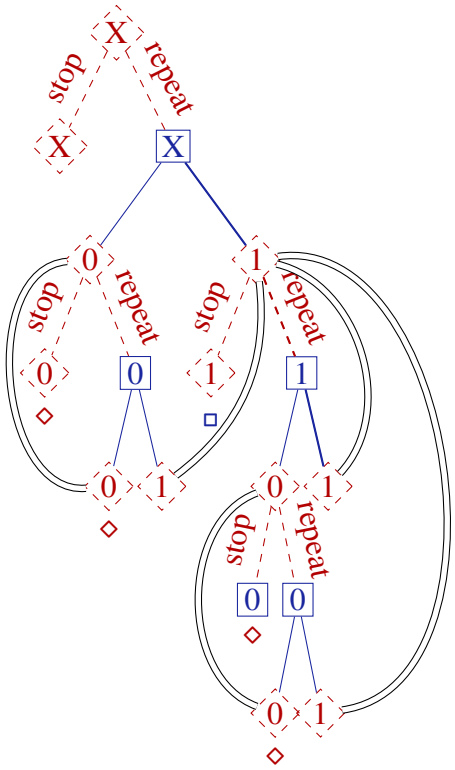

Figure 1: The filibuster game formula $\langle(x := 0 \cap x := 1)^*\rangle x = 0$ is false (unless $x = 0$ initially), but would be non-determined without least fixpoints (strategies follow thick actions). Here and in subsequent figures, we illustrate Angel's action choices by dashed edges from dashed diamonds, Demon's action choices by solid edges from solid squares, and use double lines to indicate identical states with the same continuous state and a subgame of the same structure of subsequent choices. We mark states where Angel wins by $\diamond$ and states where Demon wins by $\square$. If a winning state can be reached by a winning strategy, we enclose the mark in a circle $\textcircled{\diamond}$ or $\textcircled{\square}$, respectively.

A proof is in Appendix A. Interactive games with both duals and repetitions, however, do not generally have a Scott-continuous semantics:

$$\mathbb{R} = \varsigma_{y:=y+1^\times}(\underbrace{\bigcup_{n=1}^{\infty}(-\infty,n])}_{\mathbb{R}} \not\subseteq \bigcup_{n=1}^{\infty} \varsigma_{y:=y+1^\times}((-\infty,n]) = \emptyset$$

$$\text{since } \vDash \langle y := y+1^\times\rangle \underbrace{\exists n:\mathbb{N}\ y \leq n}_{true} \quad \text{but} \quad \nvDash \exists n:\mathbb{N}\ \langle y := y+1^\times\rangle y \leq n$$

Observe that this is related to a failure of the Barcan axiom (Section 6).

4 Operational Game Semantics

In order to relate the intuition of interactive game play to the semantics of hybrid games, we show an operational semantics for hybrid games that is more complicated than the regular modal semantics from Section 3 but makes strategies explicit and more directly reflects the intuition how hybrid games are played successively. The regular modal semantics is beneficial, because it is simpler. The operational semantics formalizes the intuition behind the game tree in Fig. 1 and relates to standard notions in game-theory. We prove in Theorem 2 below that the operational game semantics is equivalent to the regular modal semantics from Section 3. The operational semantics makes winning strategies explicit. As the set of actions A for a hybrid game, we choose:

$$\{\mathsf{l},\mathsf{r},\mathsf{s},\mathsf{g},\mathsf{d}\} \cup \{(x := \theta) : x \text{ variable}, \ \theta \text{ term}\}$$
$$\cup \{(x' = \theta \,\&\, H@r) : x \text{ variable}, \ \theta \text{ term}, H \text{ formula}, r \in \mathbb{R}_{\geq 0}\} \cup \{?\phi : \phi \text{ formula}\}$$

For game $\alpha \cup \beta$, action l decides to descend left into α, r is the action of descending right into β. In game α^*, action s decides to stop repeating, action g decides to go back and repeat. Action d starts and ends a dual game for α^d. The other actions represent assignment actions, continuous evolution actions (in which time r is the critical decision), and test actions.

The set of finite sequences of actions is denoted by $A^{(\mathbb{N})}$, the set of infinite sequences by $A^{\mathbb{N}}$. The empty sequence of actions is $()$. The concatenation, $s\hat{}t$, of sequences $s, t \in A^{(\mathbb{N})}$ is defined as $(s_1, \ldots, s_n, t_1, \ldots, t_m)$ if $s = (s_1, \ldots, s_n)$ and $t = (t_1, \ldots, t_m)$. For an $a \in A$, we write $a\hat{}t$ for $(a)\hat{}t$ and write $t\hat{}a$ for $t\hat{}(a)$. For a set $S \subseteq A^{(\mathbb{N})}$, we write $S\hat{}t$ for $\{s\hat{}t : s \in S\}$ and $t\hat{}S$ for $\{t\hat{}s : s \in S\}$. The state $\lceil t \rceil_s$ reached by playing a sequence of actions $t \in A^{(\mathbb{N})}$ from a state s is inductively defined by applying the actions sequentially, i.e., as follows:

1. $\lceil x := \theta \rceil_s = s_x^{[\![\theta]\!]_s}$
2. $\lceil x' = \theta \,\&\, H@r \rceil_s = \varphi(r)$ where $\varphi : [0, r] \to \mathcal{S}$ differentiable, $\varphi(0) = s$, $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]$ for all $\zeta \leq r$. Where $\lceil x' = \theta \,\&\, H@r \rceil_s$ is not defined if no such φ exists.
3. $\lceil ?\phi \rceil_s = \begin{cases} s & \text{if } s \in [\![\phi]\!] \\ \text{not defined} & \text{otherwise} \end{cases}$
4. $\lceil \mathsf{l} \rceil_s = \lceil \mathsf{r} \rceil_s = \lceil \mathsf{s} \rceil_s = \lceil \mathsf{g} \rceil_s = \lceil \mathsf{d} \rceil_s = \lceil () \rceil_s = s$
5. $\lceil a\hat{}t \rceil_s = \lceil t \rceil_{(\lceil a \rceil_s)}$ for $a \in A$ and $t \in A^{(\mathbb{N})}$

A *tree* is a set $T \subseteq A^{(\mathbb{N})}$ that is closed under prefixes, that is, whenever $t \in T$ and s is a prefix of t (i.e., $t = s\hat{}r$ for some $r \in A^{(\mathbb{N})}$), then $s \in T$. A node $t \in T$ is a successor of node $s \in T$ iff $t = s\hat{}a$ for some $a \in A$. By $\mathrm{leaf}(T)$ we denote the set of all leaves of T, i.e., nodes $t \in T$ that have no successor in T. The *operational game semantics* of hybrid game α is, for each state s, a tree $\mathfrak{g}(\alpha)(s) \subseteq A^{(\mathbb{N})}$ defined as follows (see Fig. 2 for a schematic illustration):

1. $\mathfrak{g}(x := \theta)(s) = \{(), (x := \theta)\}$

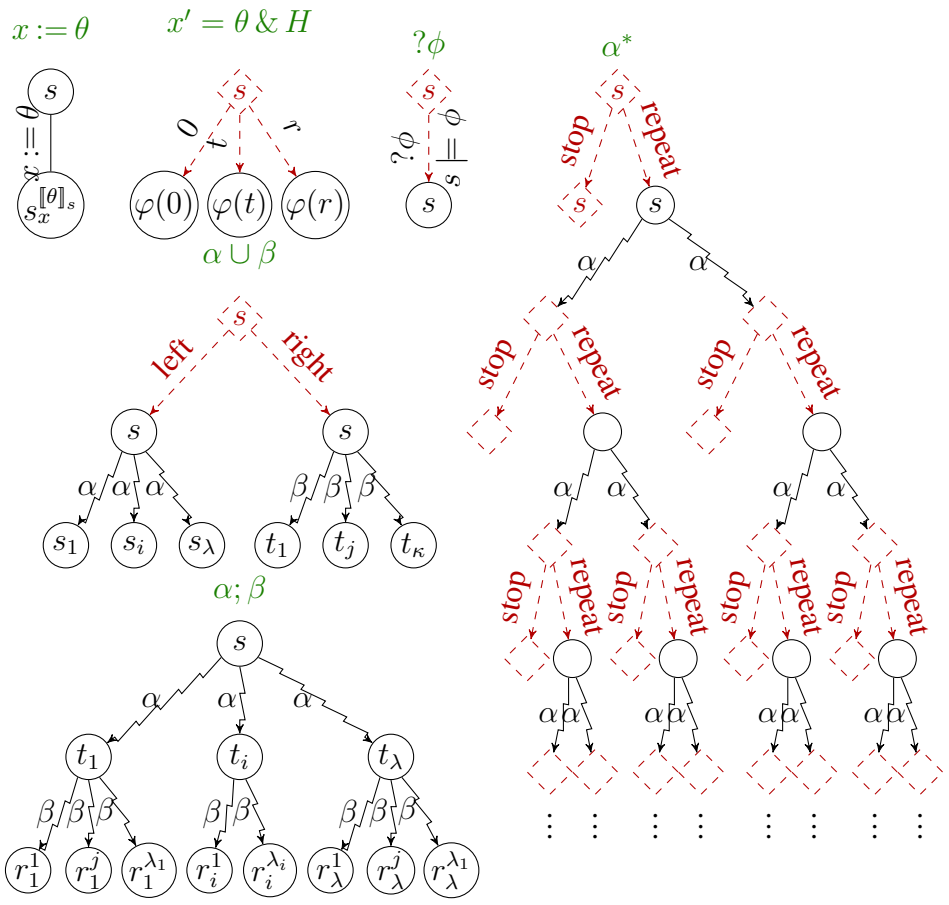

Figure 2: Operational game semantics for hybrid games of dGL

2. $\mathfrak{g}(x' = \theta \,\&\, H)(s) = \{(), (x' = \theta \,\&\, H@r) : r \in \mathbb{R}_{\geq 0}, \varphi(0) = s$ for some (differentiable) $\varphi : [0, r] \to \mathcal{S}$ such that $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]$ for all $\zeta \leq r\}$

3. $\mathfrak{g}(?\phi)(s) = \begin{cases} \{(), (?\phi)\} & \text{if } s \in [\![\phi]\!] \\ \{(), (?\mathit{false})\} & \text{otherwise} \end{cases}$

4. $\mathfrak{g}(\alpha \cup \beta)(s) = \{(), (\mathsf{l}), (\mathsf{r})\} \cup \mathsf{l}\hat{}\mathfrak{g}(\alpha)(s) \cup \mathsf{r}\hat{}\mathfrak{g}(\beta)(s)$

5. $\mathfrak{g}(\alpha; \beta)(s) = \mathfrak{g}(\alpha)(s) \cup \bigcup_{t \in \mathrm{leaf}(\mathfrak{g}(\alpha)(s))} \mathfrak{g}(\beta)(\lceil t \rceil_s)$

6. $\mathfrak{g}(\alpha^*)(s) = \bigcap \left\{ T \subseteq A^{(\mathbb{N})} : \{(), (\mathsf{s}), (\mathsf{g})\} \cup \bigcup_{t\hat{}\mathsf{g} \in \mathrm{leaf}(Z)} t\hat{}\mathsf{g}\hat{}\mathfrak{g}(\alpha)(\lceil t\hat{}\mathsf{g} \rceil_s)\hat{}\{(), (\mathsf{s}), (\mathsf{g})\} \subseteq T \right\}.$

7. $\mathfrak{g}(\alpha^d)(s) = \{(), (\mathsf{d})\} \cup \mathsf{d}\hat{}\mathfrak{g}(\alpha)(s)\hat{}\{(), \mathsf{d}\}$

Angel gets to choose which action to take at node $t \in \mathfrak{g}(\alpha)(s)$ if t has an even number of occurrences of d, otherwise Demon gets to choose. In the former case we say *Angel acts at* t, in the

latter *Demon acts at* t. If the player who chooses the action at $t \in \mathfrak{g}(\alpha)(s)$ is deadlocked, because the only successor actions have a condition that is not satisfied like $?\mathit{false}$ or $x' = \theta \,\&\, x \geq 0$ at a state where $x < 0$, then that player loses immediately.

A *strategy for Angel* from initial state s is a nonempty subtree $\sigma \subseteq \mathfrak{g}(\alpha)(s)$ such that

1. for all $t \in \sigma$ at which Demon acts, $t\hat{}a \in \sigma$ for all $a \in A$ such that $t\hat{}a \in \mathfrak{g}(\alpha)(s)$.
2. for all $t \in \sigma$ at which Angel acts, $t \notin \mathrm{leaf}(\mathfrak{g}(\alpha)(s))$, there is a unique $a \in A$ with $t\hat{}a \in \sigma$.

Strategies for Demon are defined accordingly, with "Angel" and "Demon" swapped. The action sequence $\sigma \oplus \tau$ played from state s when Angel plays strategy σ and Demon plays strategy τ from s is defined as the sequence $(a_1, \ldots, a_n) \in A^{(\mathbb{N})}$ of maximal length such that

$$a_{n+1} := \begin{cases} a & \text{if Angel acts at } (a_1, \ldots, a_n) \text{ and } (a_1, \ldots, a_n)\hat{}a \in \sigma \\ a & \text{if Demon acts at } (a_1, \ldots, a_n) \text{ and } (a_1, \ldots, a_n)\hat{}a \in \tau \\ \text{not defined} & \text{otherwise} \end{cases}$$

By definition of a strategy for Angel/Demon, the a is unique. A *winning strategy for Angel* for winning condition $X \subseteq \mathcal{S}$ from state s is a strategy $\sigma \subseteq \mathfrak{g}(\alpha)(s)$ for Angel from s such that, for all strategies $\tau \subseteq \mathfrak{g}(\alpha)(s)$ for Demon from s: Demon deadlocks or $\lceil \sigma \oplus \tau \rceil_s \in X$. A *winning strategy for Demon* for (Demon's) winning condition $X \subseteq \mathcal{S}$ from state s is a strategy $\tau \subseteq \mathfrak{g}(\alpha)(s)$ for Demon from s such that, for all strategies $\sigma \subseteq \mathfrak{g}(\alpha)(s)$ for Angel from s: Angel deadlocks or $\lceil \sigma \oplus \tau \rceil_s \in X$. By definition, it cannot be that Angel has a winning strategy for X from s and, at the same time, Demon has a winning strategy for $\mathcal{S} \setminus X$ from s. If we understand $[\alpha]\phi$ as Demon having a strategy to achieve ϕ, this justifies the *consistency* direction $\neg(\langle\alpha\rangle\phi \wedge [\alpha]\neg\phi)$ of $[\alpha]\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi$. Determinacy, i.e., the direction $\langle\alpha\rangle\phi \vee [\alpha]\neg\phi$ of $[\alpha]\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi$, holds by definition in the regular modal semantics of Section 3, but can now be justified in the operational semantics based on the Borel determinacy theorem [Mar75].

Theorem 1 (Determinacy)**.** *Hybrid games are* determined*, i.e., for any hybrid game α, initial state s, and winning condition $X \subseteq \mathcal{S}$, either Angel has a winning strategy for X from s or Demon has a winning strategy for $\mathcal{S} \setminus X$ from s.*

A proof is in Appendix B. We show that the regular modal semantics from Section 3 is equivalent to the operational semantics (proof in Appendix C):

Theorem 2 (Equivalent semantics)**.** *The regular modal semantics of* dGL *is equivalent to the game tree operational semantics of* dGL*, i.e., for each hybrid game α, each initial state s, and each winning condition $X \subseteq \mathcal{S}$ for Angel:*

$$s \in \varsigma_\alpha(X) \iff \textit{there is a winning strategy } \sigma \subseteq \mathfrak{g}(\alpha)(s) \textit{ for Angel for } X \textit{ from } s$$

5 Proof Calculus

Simple dGL formulas can be checked by a simple tableau procedure that expands the options of all players and detects loops for termination as shown in the game tree examples. This does not extend to more general dGL formulas, however, which have inherently infinite states. In Fig. 3, we present a proof calculus for proving validity of general dGL formulas.

$$\langle{:=}\rangle\ \langle x := \theta\rangle\phi(x) \leftrightarrow \phi(\theta)$$

$$\langle ?\rangle\ \langle ?\psi\rangle\phi \leftrightarrow (\psi \wedge \phi)$$

$$\langle'\rangle\ \langle x' = \theta\rangle\phi \leftrightarrow \exists t{\geq}0\, \langle x := y(t)\rangle\phi \qquad (y'(t) = \theta)$$

$$\langle\cup\rangle\ \langle \alpha \cup \beta\rangle\phi \leftrightarrow \langle\alpha\rangle\phi \vee \langle\beta\rangle\phi$$

$$\langle;\rangle\ \langle\alpha;\beta\rangle\phi \leftrightarrow \langle\alpha\rangle\langle\beta\rangle\phi$$

$$\langle^*\rangle\ \phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi \rightarrow \langle\alpha^*\rangle\phi$$

$$\langle^d\rangle\ \langle\alpha^d\rangle\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi$$

$$\overleftarrow{\text{B}}\ \exists x\, \langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\exists x\, \phi \qquad (x \notin \alpha)$$

$$\text{R}\ \frac{\phi \rightarrow \psi}{\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\psi}$$

$$\text{FP}\ \frac{\phi \vee \langle\alpha\rangle\psi \rightarrow \psi}{\langle\alpha^*\rangle\phi \rightarrow \psi}$$

$$\text{con}\ \frac{\varphi(v) \wedge v > 0 \rightarrow \langle\alpha\rangle\varphi(v-1)}{\varphi(v) \rightarrow \langle\alpha^*\rangle\exists v{\leq}0\, \varphi(v)} \qquad (v \notin \alpha)$$

Figure 3: Differential game logic proof rules

The proof calculus of dGL shares several axioms with the proof calculus of dL [Pla08, Pla12b, Pla12a]. We use the first-order Hilbert calculus (modus ponens and $\forall$-generalization) as a basis and allow all instances of valid formulas of first-order real arithmetic as axioms, which are decidable [Tar51]. Axiom $\langle{:=}\rangle$ is Hoare's assignment rule. Formula $\phi(\theta)$ is obtained from $\phi(x)$ by *substituting* θ for x, provided x does not occur in the scope of a quantifier or modality binding x or a variable of θ. A modality $\langle\alpha\rangle$ containing $z := $ or z' *binds* z. In axiom $\langle'\rangle$, $y(\cdot)$ is the (unique [Wal98, Theorem 10.VI]) solution of the symbolic initial value problem $y'(t) = \theta, y(0) = x$. It goes without saying that variables like t are fresh in Fig. 3. Axioms $\langle ?\rangle$, $\langle\cup\rangle$, and $\langle;\rangle$ are as in dL [Pla12b]. Axiom $\langle^*\rangle$ is the iteration axiom. The converse of $\langle^*\rangle$ can be derived[1] and is also denoted by $\langle^*\rangle$. Axiom $\langle^d\rangle$ is specific to dGL and characterizes dual games. Recall $\neg\langle\alpha\rangle\neg\phi \equiv [\alpha]\phi$. Axiom $\langle^d\rangle$ says that Angel has a winning strategy for ϕ in dual game α^d iff Demon has a winning strategy for ϕ in α. Axiom $\overleftarrow{\text{B}}$ is the converse Barcan formula of first-order modal logic, characterizing monotonic domains [HC96]. In order for it to be sound for dGL, x must not occur in α, written $x \notin \alpha$.

Rule R is the generalization rule of regular modal logic C. Rule FP is the fixpoint rule, characterizing $\langle\alpha^*\rangle\phi$ as a smallest fixpoint. Rule con, in which v does not occur in α, is a variation of Harel's convergence rule, suitably adapted to hybrid games over $\mathbb{R}$. It expresses that, if Angel has a strategy to make progress from $\varphi(v)$ to $\varphi(v-1)$ along α, then, from any state where $\varphi(v)$ holds, she has a strategy to reach $\varphi(v)$ for some $v \leq 0$ by repeating α.

[1] $\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi \rightarrow \langle\alpha^*\rangle\phi$ is valid by $\langle^*\rangle$. Thus, $\langle\alpha\rangle(\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi) \rightarrow \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by R. Hence, $\phi \vee \langle\alpha\rangle(\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi) \rightarrow \phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by propositional congruence. Consequently, $\langle\alpha^*\rangle\phi \rightarrow \phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by FP.

Example 1. The dual filibuster game formula from Section 3 can be proved as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{*}{x=0\to 0=0\vee 1=0}\mathbb{R}}{x=0\to\langle x:=0\rangle x=0\vee\langle x:=1\rangle x=0}\langle:=\rangle}{x=0\to\langle x:=0\cup x:=1\rangle x=0}\langle\cup\rangle}{x=0\to[x:=0\cap x:=1]x=0}\langle^d\rangle}{x=0\to[(x:=0\cap x:=1)^*]x=0}\text{ind}}{x=0\to\langle(x:=0\cup x:=1)^\times\rangle x=0}\langle^d\rangle$$

Almost the same dGL proof proves $x=0\to\langle(x:=x\cup x:=1)^\times\rangle x=0$. We note that significantly more challenging systems with complex hybrid dynamics are provable in the dGL calculus.

The primary difference of the axiomatization of dGL compared to differential dynamic logic [Pla12a] is the addition of axiom $\langle^d\rangle$ for dual games, the absence of axiom K, absence of the Barcan formula (dGL only has the converse Barcan axiom $\overleftarrow{\text{B}}$), and absence of Gödel's necessitation rule (dGL only has the regular modal rule R). Given the big semantical difference of run versus game, it is striking to see this concise difference in axioms. This indicates that we have found the right logical characterizations. Due to the absence of K, we will see (in Section 6) why the induction axiom and the convergence axiom are also absent in dGL, while corresponding rules are still valid. The induction rule (ind, which is derivable from FP) and the convergence rule (con) are sound for dGL (a proof is in Appendix D).

Lemma 2. *Rule FP and the induction rule (ind) of dynamic logic are interderivable in the* dGL *calculus:*

$$\textit{(ind)}\quad \frac{\psi\to[\alpha]\psi}{\psi\to[\alpha^*]\psi}$$

Theorem 3 (Soundness)**.** *The* dGL *proof rules in Fig. 3 are sound.*

A proof is in Appendix D. The proof rules in Fig. 3 do not handle differential equations with evolution domain constraints (other than $true$). Unlike in (poor test) differential dynamic logic [Pla08, Pla10, Pla12a], however, every hybrid game containing a differential equation with evolution domain constraints can be replaced equivalently by a hybrid game without evolution domain constrains (even with poor tests, i.e., each test $?\phi$ uses only first-order formulas ϕ)!

Lemma 3. *Evolution domains of differential equations are definable as hybrid games. That is, for every hybrid game α, there is a hybrid game β that is equivalent (i.e., $\varsigma_\alpha(X)=\varsigma_\beta(X)$ for all X) but has no evolution domain constraints.*

Proof. When, for notational convenience, we assume the (vectorial) differential equation $x'=\theta(x)$ to contain a clock $x_0'=1$ and that t_0 and z are fresh variables, then the following two hybrid games are equivalent:

$$x'=\theta(x)\,\&\,H(x)\ \equiv\ t_0:=x_0;x'=\theta(x);(z:=x;z'=-\theta(z))^d;?(z_0\geq t_0\to H(z)) \qquad (1)$$

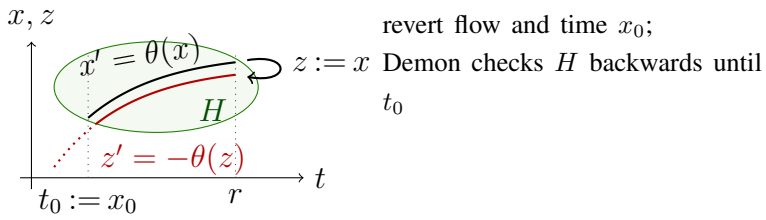

Figure 4: Angel evolves x forwards in time along $x' = \theta(x)$, Demon checks evolution domain backwards in time along $z' = -\theta(z)$ on a copy z of the state

See Fig. 4 for an illustration. Suppose the current player is Angel. The idea behind game equivalence (1) is that the fresh variable t_0 remembers the initial time x_0, and Angel then evolves along $x' = \theta(x)$ for any amount of time (Angel's choice). Afterwards, the opponent Demon copies the state x into a fresh variable (vector) z that it can evolve backwards along $(z' = -\theta(z))^d$ for any amount of time (Demon's choice). The original player Angel must then pass the challenge $?(z_0 \geq t_0 \to H(z))$, i.e., Angel loses immediately if Demon was able to evolve backwards and leave region $H(z)$ while satisfying $z_0 \geq t_0$, which checks that Demon did not evolve backward for longer than Angel evolved forward. Otherwise, when Angel passes the test, the extra variables t_0, z become irrelevant (they are fresh) and the game continues from the current state x that Angel chose in the first place (by selecting a duration for the evolution that Demon could not invalidate). □

6 Separating Axioms

In order to illustrate how and why dGL differs from differential dynamic logic dL [Pla08, Pla12a], i.e., how reasoning about hybrid games really differs from reasoning about hybrid systems, we identify separating axioms, that is, axioms of dL that do not hold in dGL. For each such fundamental separating axiom, we give a simple counterexample illustrating what makes the hybrid game focus of dGL behave differently than hybrid systems. First, we show that dGL only is a regular modal logic, while dL is a normal modal logic [HC96]. Axiom K, the modal modus ponens from modal logic [HC96], dynamic logic [Pra76], and differential dynamic logic [Pla12a]:

$$[\alpha](\phi \to \psi) \to ([\alpha]\phi \to [\alpha]\psi)$$

is not sound for dGL as witnessed using the choice $\alpha \equiv (x := 1 \cap x := 0); y := 0$ and $\phi \equiv x = 1$, $\psi \equiv y = 1$; see Fig. 5. The global rule version of K, i.e., the implicative version of Gödel's generalization rule is still sound and derives with $\langle^d\rangle$ from R using $\alpha \equiv \beta^d$

$$\frac{\phi \to \psi}{[\beta]\phi \to [\beta]\psi}$$

The normal Gödel generalization rule G, i.e.,

$$\frac{\phi}{[\alpha]\phi}$$

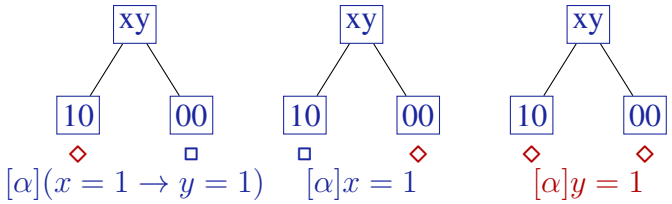

Figure 5: Game trees for counterexample to axiom K using $\alpha \equiv (x := 1 \cap x := 0); y := 0$.

however, is not sound for dGL as witnessed by the choice $\alpha \equiv (?false)^d, \phi \equiv true$.

The Barcan axiom B, which characterizes anti-monotonic domains in first-order modal logic [HC96], is sound for constant-domain first-order dynamic logic and for differential dynamic logic dL when x does not occur in α [Pla12a]

$$\langle\alpha\rangle\exists x\, \phi \rightarrow \exists x\, \langle\alpha\rangle\phi \qquad (x \notin \alpha)$$

but, unlike the converse Barcan $\overleftarrow{\text{B}}$, the Barcan axiom is not sound for dGL as witnessed by the choice $\alpha \equiv y := y + 1^\times$ and $\phi \equiv x \geq y$. The equivalent Barcan formula

$$\forall x\, [\alpha]\phi \rightarrow [\alpha]\forall x\, \phi \qquad (x \notin \alpha)$$

is not sound for dGL as witnessed by the choice $\alpha \equiv y := y + 1^\times$ and $\phi \equiv y \geq x$.

The first arrival axiom, $\langle\alpha^*\rangle\phi \rightarrow \phi \vee \langle\alpha^*\rangle(\neg\phi \wedge \langle\alpha\rangle\phi)$, which holds for dL, expresses that, if ϕ holds after a repetition of α, then it either holds right away or α can be repeated so that ϕ does not hold yet but can hold after one more repetition. This axiom does not hold, however, for dGL as witnessed by $\alpha^* \equiv ((x := x - y \cap x := 0); y := x)^*$ and $\phi \equiv x = 0$, since two iterations surely yield $x = 0$, but one iteration may or may not yield $x = 0$, depending on Demon's choice; see Fig. 6.

Unlike induction rule ind, induction axiom $[\alpha^*](\phi \rightarrow [\alpha]\phi) \rightarrow (\phi \rightarrow [\alpha^*]\phi)$, which is the dual of the first arrival axiom, holds for dL, but does not hold for dGL as witnessed by

$$\alpha^* \equiv ((x := a; a := 0) \cap x := 0)^*$$

and $\phi \equiv x = 1$; see Fig. 7.

Note that the failure of the induction axiom in this counterexample hinges on the fact that Angel is free to decide whether or not to repeat α after each round depending on the state. This would be different if we had chosen an *advance notice semantics* for α^* in which the number of times that game α will be repeated would have to be announced by the player when the loop begins. In this example, if Angel announces that she has chosen n repetitions of the game, then Demon wins (for $a \neq 0$) by choosing the $x := 0$ option $n - 1$ times followed by one choice of $x := a; a := 0$. Such games that need a prior commitment from the player on the number of repetitions before α^* starts would lead to a very different semantics. If we had chosen an advance notice semantics, then the following formula would be valid, but it is not valid in dGL (see Fig. 7 right):

$$x = 1 \wedge a = 1 \rightarrow [((x := a; a := 0) \cap x := 0)^*]x = 1 \tag{2}$$

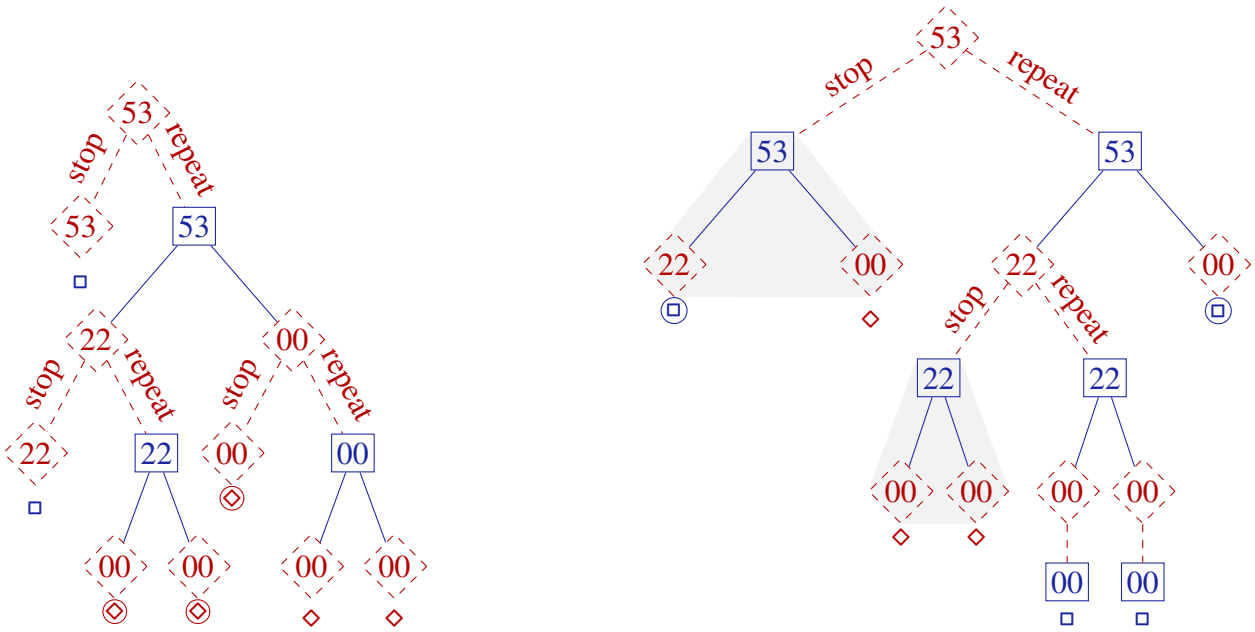

Figure 6: Game trees for counterexample to first arrival axiom (notation: x, y) with game $\alpha \equiv (x := x - y \cap x := 0); y := x$. **(left)** $\langle \alpha^* \rangle x = 0$ is true no matter which choices Demon makes **(right)** $\langle \alpha^* \rangle (x \neq 0 \wedge \langle \alpha \rangle x = 0)$ is false, because stop can be defeated by $x := x - y$ and repeat can be defeated by $x := 0$.

The dual formula, instead, is valid in dGℒ but not with advance notice (Fig. 7):

$$x = 1 \wedge a = 1 \rightarrow \langle ((x := a; a := 0) \cap x := 0)^* \rangle x \neq 1$$

Our semantics is more general, because advance notice games can be expressed easily in dGℒ by having the players choose a counter c before the loop that decreases to 0 during the repetition. The advance notice semantics for (2) can be expressed in dGℒ, e.g., as

$$x = 1 \wedge a = 1 \rightarrow [c := 0; c := c + 1^*; (((x := a; a := 0) \cap x := 0); c := c - 1)^*; ?c = 0]x = 1$$

The dGℒ semantics cannot, however, be expressed conversely in an advance notice semantics, so the dGℒ semantics is strictly more general. Many other game interactions can be defined in similar ways from the elementary operators that dGℒ provides.

7 Related Work

Discrete games and the interaction of games and logic for various purposes have been studied with much success [vNM55, Par85, HS97, PP03]. Differential games have been studied exhaustively [Isa67] and are of interest to understand phenomena in games where the actions are continuous in time. Here, we look at the complementary model of hybrid games where the underlying system

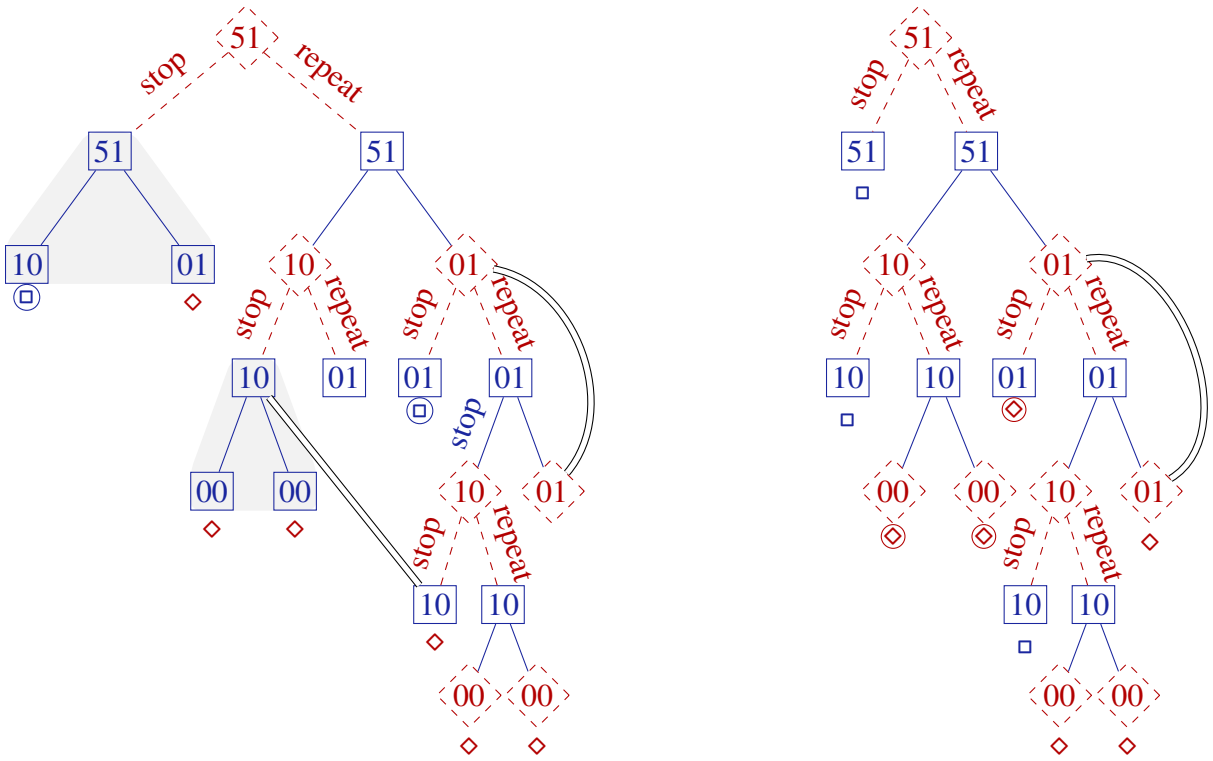

Figure 7: Game trees for counterexample to induction axiom (notation: x, a) with game $\alpha \equiv (x := a; a := 0) \cap x := 0$. **(left)** $[\alpha^*](x = 1 \rightarrow [\alpha]x = 1)$ is true by the strategy "if Angel chose stop, choose $x := a; a := 0$, otherwise always choose $x := 0$" **(right)** $[\alpha^*]x = 1$ is false by strategy "repeat once and repeat once more if $x = 1$, then stop"

is that of a hybrid system with interacting discrete and continuous dynamics, but the game actions are chosen at discrete instants of time, even if they take effect in continuous time.

Reachability aspects of games for hybrid systems have been studied before. A game view on hybrid systems verification has been proposed by Tomlin and coauthors following a Hamilton-Jacobi-Bellman PDE formulation [TMBO03, MBT05], with subsequent extensions by Gao et al. [GLQ07]. Their primary focus is on adversarial choices in the continuous dynamics, which is very interesting, but not what we consider here. It is also easier to get the axioms of our proof calculus sound than numerical approximations of PDEs. WCTL properties of STORMED hybrid games, which require monotonicity properties for the system evolution, have been shown to be decidable using bisimulation quotients [VPVD11]. The special case of o-minimal hybrid games has been shown to be decidable earlier by Bouyer et al. [BBC07]. The case of rectangular hybrid games is known to be decidable [HHM99].

We take a complementary view and study logics and proofs for hybrid games instead of searching for decidable fragments using bisimulation quotients [HHM99, BBC07, VPVD11]. Our notion of hybrid games has more flexible nested hybrid choices for the agents than the fixed controller-plant interaction considered in related work. We consider more general logical formulas.

8 Conclusions and Future Work

We have presented *differential game logic* (dGL) for hybrid games, which unifies differential dynamic logic (dL) and Parikh's game logic. We have provided a regular modal semantics for dGL, a proof calculus, and proved soundness. Our logical setting enables us to characterize the essential logical difference of hybrid systems proving compared to hybrid games proving by identifying the axioms that separate dL and dGL: the axiom of duality, axiom K, Barcan axiom, and Gödel's generalization rule (replaced with the regular rule). We observe that there is a striking similarity of our dGL proof calculus with our calculus for stochastic differential dynamic logic SdL [Pla11, Pla12b], despite their fundamentally different semantical presuppositions (adversarial nondeterminism versus stochasticity). This leads us to conjecture the existence of a deeper logical connection relating stochastic and adversarial uncertainty.

Acknowledgments

I want to thank Erik Zawadzki for helpful discussions about game theory.

References

[AM98] Eugene Asarin and Oded Maler. Achilles and the tortoise climbing up the arithmetical hierarchy. *J. Comput. Syst. Sci.*, 57(3):389–398, 1998.

[BBC07] Patricia Bouyer, Thomas Brihaye, and Fabrice Chevalier. Weighted o-minimal hybrid systems are more decidable than weighted timed automata! In Sergei N. Artëmov and Anil Nerode, editors, *LFCS*, volume 4514 of *LNCS*, pages 69–83. Springer, 2007.

[CL00] Franck Cassez and Kim Guldstrand Larsen. The impressive power of stopwatches. In *CONCUR*, pages 138–152, 2000.

[DBL12] *Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2012, June 2528, 2012, Dubrovnik, Croatia*, 2012.

[DN00] Jennifer M. Davoren and Anil Nerode. Logics for hybrid systems. *IEEE*, 88(7):985–1010, 2000.

[GLQ07] Y. Gao, J. Lygeros, and M. Quincampoix. On the reachability problem for uncertain hybrid systems. *IEEE T. Automat. Contr.*, 52(9):1572–1586, September 2007.

[HC96] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, 1996.

[Hen96] Thomas A. Henzinger. The theory of hybrid automata. In *LICS*, pages 278–292, 1996.

[HHM99] Thomas A. Henzinger, Benjamin Horowitz, and Rupak Majumdar. Rectangular hybrid games. In Jos C. M. Baeten and Sjouke Mauw, editors, *CONCUR*, volume 1664 of *LNCS*, pages 320–335. Springer, 1999.

[HS97] Jaakko Hintikka and Gabriel Sandu. Game-theoretical semantics. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*. Elsevier, 1997.

[Isa67] Rufus Philip Isaacs. *Differential Games*. John Wiley, 1967.

[Kec94] Alexander S. Kechris. *Classical Descriptive Set Theory*. Springer, 1994.

[Koz06] Dexter Kozen. *Theory of Computation*. Springer, 2006.

[Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

[MBT05] Ian Mitchell, Alexandre M. Bayen, and Claire Tomlin. A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE T. Automat. Contr.*, 50(7):947–957, 2005.

[Nas51] John Nash. Non-cooperative games. *Ann. Math.*, 54(2):286–295, 1951.

[Par85] Rohit Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–140, 1985. M. Karpinski and J. van Leeuwen, eds., Topics in the Theory of Computation.

[Pla08] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008.

[Pla10] André Platzer. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg, 2010.

[Pla11] André Platzer. Stochastic differential dynamic logic for stochastic hybrid programs. In Nikolaj Bjørner and Viorica Sofronie-Stokkermans, editors, *CADE*, volume 6803 of *LNCS*, pages 431–445. Springer, 2011.

[Pla12a] André Platzer. The complete proof theory of hybrid systems. In *LICS* [DBL12].

[Pla12b] André Platzer. Logics of dynamical systems (invited tutorial). In *LICS* [DBL12].

[PP03] Marc Pauly and Rohit Parikh. Game logic - an overview. *Studia Logica*, 75(2):165–182, 2003.

[Pra76] Vaughan R. Pratt. Semantical considerations on Floyd-Hoare logic. In *FOCS*, 1976.

[Tar51] Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, Berkeley, 2nd edition, 1951.

[TLS00] Claire J. Tomlin, John Lygeros, and Shankar Sastry. A game theoretic approach to controller design for hybrid systems. *Proc. IEEE*, 88(7):949–970, 2000.

[TMBO03] Claire Tomlin, Ian Mitchell, Alexandre M. Bayen, and Meeko Oishi. Computational techniques for the verification of hybrid systems. *Proc. IEEE*, 91(7):986–1001, 2003.

[TPS98] Claire Tomlin, George J. Pappas, and Shankar Sastry. Conflict resolution for air traffic management: a study in multi-agent hybrid systems. *IEEE T. Automat. Contr.*, 43(4):509–521, 1998.

[vNM55] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton Univ. Press, 3rd edition, 1955.

[VPVD11] Vladimeros Vladimerou, Pavithra Prabhakar, Mahesh Viswanathan, and Geir E. Dullerud. Specifications for decidable hybrid games. *Theor. Comput. Sci.*, 412(48):6770–6785, 2011.

[Wal98] Wolfgang Walter. *Ordinary Differential Equations*. Springer, 1998.

A Proof of Scott-Continuity

We provide a proof of the result about Scott-continuity.

Proof of Lemma 1. By monotonicity, $\bigcup_{n\in I} \varsigma_\alpha(X_n) \subseteq \varsigma_\alpha(\bigcup_{n\in I} X_n)$. We show the converse inclusion by induction on the structure of α: $\varsigma_\alpha(\bigcup_{n\in I} X_n) \subseteq \bigcup_{n\in I} \varsigma_\alpha(X_n)$.

1. $\varsigma_{x:=\theta}(\bigcup_{n\in I} X_n) = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in \bigcup_{n\in I} X_n\} \subseteq \bigcup_{n\in I}\{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in X_n\} = \bigcup_{n\in I} \varsigma_{x:=\theta}(X_n)$, because $s_x^{[\![\theta]\!]_s} \in \bigcup_{n\in I} X_n$ implies $s_x^{[\![\theta]\!]_s} \in X_n$ for some n.

2. $\varsigma_{x'=\theta\,\&\,H}(\bigcup_{n\in I} X_n) = \{\varphi(0) \in \mathcal{S} : \frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]$ for all $\zeta \leq r$ for some (differentiable) $\varphi : [0, r] \to \mathcal{S}$ such that $\varphi(r) \in \bigcup_{n\in I} X_n\} \subseteq \bigcup_{n\in I} \varsigma_{x'=\theta\,\&\,H}(X_n) = \{\varphi(0) \in \mathcal{S} : \ldots \varphi(r) \in X_n\}$, because $\varphi(r) \in \bigcup_{n\in I} X_n$ implies $\varphi(r) \in X_n$ for some n.

3. $\varsigma_{?\phi}(\bigcup_{n\in I} X_n) = [\![\phi]\!] \cap \bigcup_{n\in I} X_n = \bigcup_{n\in I}([\![\phi]\!] \cap X_n) = \bigcup_{n\in I} \varsigma_{?\phi}(X_n)$

4. $\varsigma_{\alpha\cup\beta}(\bigcup_{n\in I} X_n) = \varsigma_\alpha(\bigcup_{n\in I} X_n) \cup \varsigma_\beta(\bigcup_{n\in I} X_n) \overset{\text{IH}}{=} (\bigcup_{n\in I} \varsigma_\alpha(X_n)) \cup (\bigcup_{n\in I} \varsigma_\beta(X_n))$
$= \bigcup_{n\in I}(\varsigma_\alpha(X_n) \cup \varsigma_\beta(X_n)) = \bigcup_{n\in I} \varsigma_{\alpha\cup\beta}(X_n)$

5. $\varsigma_{\alpha;\beta}(\bigcup_{n\in I} X_n) = \varsigma_\alpha(\varsigma_\beta(\bigcup_{n\in I} X_n)) \overset{\text{IH}}{=} \varsigma_\alpha(\bigcup_{n\in I} \varsigma_\beta(X_n)) \overset{\text{IH}}{=} \bigcup_{n\in I} \varsigma_\alpha(\varsigma_\beta(X_n)) = \bigcup_{n\in I} \varsigma_{\alpha;\beta}(X_n)$

6. $\varsigma_{\alpha^*}(\bigcup_{n\in I} X_n) = \mu Z.(\bigcup_{n\in I} X_n) \cup \varsigma_\alpha(Z) = (\bigcup_{n\in I} X_n) \cup \varsigma_\alpha(\varsigma_{\alpha^*}(\bigcup_{n\in I} X_n))$ is the least fixpoint. We will show that $\bigcup_{n\in I} \varsigma_{\alpha^*}(X_n)$ also is a fixpoint, implying $\varsigma_{\alpha^*}(\bigcup_{n\in I} X_n) \subseteq \bigcup_{n\in I} \varsigma_{\alpha^*}(X_n)$. Indeed, $(\bigcup_{n\in I} X_n) \cup \varsigma_\alpha(\bigcup_{n\in I} \varsigma_{\alpha^*}(X_n)) \overset{\text{IH}}{=} (\bigcup_{n\in I} X_n) \cup \bigcup_{n\in I} \varsigma_\alpha(\varsigma_{\alpha^*}(X_n)) = \bigcup_{n\in I}(X_n \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X_n)) \overset{\text{fix}}{=} \bigcup_{n\in I} \varsigma_{\alpha^*}(X_n)$. □

B Determinacy Proof

In this section, we prove determinacy (Theorem 1) using the operational semantics of dGL based on the Borel determinacy theorem.

Theorem 4 (Borel determinacy theorem [Mar75, Kec94, Theorem 20.6])**.** *Let T a nonempty pruned tree on a A and let $X \subseteq [T]$ Borel in the product topology on $A^{\mathbb{N}}$ induced by the discrete topology on A. Then the Gale-Stewart game with rules T and winning condition X is determined.*

With this deep result from the literature, we can prove determinism of dGL (Theorem 1):

Proof of Theorem 1. Determinacy follows from the Borel determinacy theorem (Theorem 4), because there are no draws and all plays have (unbounded) finite length since Angel and Demon, respectively, can only choose to repeat α^* and $\alpha^\times$, respectively, finitely often (repetition is defined by a least fixpoint). For this we show that the winning condition is open in the product topology on the action sequences $A^{\mathbb{N}}$ induced by the discrete topology on the action set A. To see this, note that the set of those sequences is a union of sets of the form $\{t\hat{}r : r \in A^{\mathbb{N}}\}$ for some finite action sequence $t \in A^{(\mathbb{N})}$, which are open in the product topology on $A^{\mathbb{N}}$. Furthermore, arbitrary unions

of open sets are open. In particular, the winning conditions are Borel in the product topology on $A^{\mathbb{N}}$ induced by the discrete topology on A.[2] A hybrid game can be cast easily as a Gale-Stewart game (a game of infinite length in which the players alternate strictly), which is assumed by the Borel determinacy theorem, just by adding a stuttering action f to A. The stuttering action defined by $\lceil \mathsf{f} \rceil_s = s$ is the only action that a player can choose in the Gale-Stewart game when the next move in the hybrid game $\mathfrak{g}(\alpha)(s)$ is not his choice or the hybrid game has terminated already. □

C Equivalence of Regular Modal and Operational Semantics

We prove equivalence of the regular modal semantics from Section 3 and the operational game semantics from Section 4.

Proof of Theorem 2. We proceed by induction on the structure of α (and, simultaneously, on the number of times repetitions in α are repeated) and prove equivalence. As part of the equivalence proof, we construct a winning strategy σ achieving X using that $s \in \varsigma_\alpha(X)$.

1. $s \in \varsigma_{x:=\theta}(X) \iff s_x^{[\![\theta]\!]_s} \in X \iff \lceil \sigma \oplus \tau \rceil_s = \lceil x := \theta \rceil_s = s_x^{[\![\theta]\!]_s} \in X$, using $\sigma \overset{\text{def}}{=} \{(x := \theta)\}$.

2. $s \in \varsigma_{x'=\theta \,\&\, H}(X) \iff s = \varphi(0), \varphi(r) \in X$ for some $r \in \mathbb{R}$ and some (differentiable) $\varphi : [0, r] \to \mathcal{S}$ such that $\frac{\mathsf{d}\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]$ for all $\zeta \leq r \iff \lceil \sigma \oplus \tau \rceil_s = \lceil x' = \theta \,\&\, H@r \rceil_s = \varphi(r) \in X$, using $\sigma \overset{\text{def}}{=} \{(x' = \theta \,\&\, H@r)\}$.

3. $s \in \varsigma_{?\phi}(X) = [\![\phi]\!] \cap X \iff \lceil \sigma \oplus \tau \rceil_s = \lceil ?\phi \rceil_s = s \in X$, using $\sigma \overset{\text{def}}{=} \{(?\phi)\}$.

4. $s \in \varsigma_{\alpha \cup \beta}(X) = \varsigma_\alpha(X) \cup \varsigma_\beta(X) \iff s \in \varsigma_\alpha(X)$ or $s \in \varsigma_\beta(X)$. By induction hypothesis, this is equivalent to: there is a winning strategy $\sigma_\alpha \subseteq \mathfrak{g}(\alpha)(s)$ for Angel for X from s or there is a winning strategy $\sigma_\beta \subseteq \mathfrak{g}(\beta)(s)$ for Angel for X from s. This is equivalent to $\sigma \subseteq \mathfrak{g}(\alpha \cup \beta)(s)$ being a winning strategy for Angel for X from s, using either $\sigma \overset{\text{def}}{=} \{(\mathsf{l})\} \cup \mathsf{l}\hat{}\sigma_\alpha$ or $\sigma \overset{\text{def}}{=} \{(\mathsf{r})\} \cup \mathsf{r}\hat{}\sigma_\beta$.

5. $s \in \varsigma_{\alpha;\beta}(X) = \varsigma_\alpha(\varsigma_\beta(X))$ By induction hypothesis, this is equivalent to the existence of a strategy $\sigma_\alpha \subseteq \mathfrak{g}(\alpha)(s)$ for Angel such that for all strategies $\tau \subseteq \mathfrak{g}(\alpha)(s)$ for Demon: $\lceil \sigma_\alpha \oplus \tau \rceil_s \in \varsigma_\beta(X)$. By induction hypothesis, $\lceil \sigma_\alpha \oplus \tau \rceil_s \in \varsigma_\beta(X)$ is equivalent to the existence of a winning strategy σ_τ for Angel (which depends on the state $\lceil \sigma_\alpha \oplus \tau \rceil_s$ that the previous α game led to) with winning condition X from $\lceil \sigma_\alpha \oplus \tau \rceil_s$. This is equivalent to $\sigma \subseteq \mathfrak{g}(\alpha;\beta)(s)$ being a winning strategy for Angel for X from s, using

$$\sigma \overset{\text{def}}{=} \sigma_\alpha \cup \bigcup_{\sigma_\alpha \oplus \tau} (\sigma_\alpha \oplus \tau)\hat{}\sigma_\tau$$

[2] Observe that the winning conditions are Borel in a different topology than the Euclidean topology.

The union is over all leaves $\sigma_\alpha \oplus \tau$ for which the game is not won by a player yet. Note that σ is a winning strategy for X, because, for all plays for which the game is decided during α, the strategy σ_α already wins the game. For the others, σ_τ wins the game from the respective state $\lceil \sigma_\alpha \oplus \tau \rceil_s$ that was reached by the actions $\sigma_\alpha \oplus \tau$ according to Demon's strategy τ.

6. We prove the case α^* using a simultaneous induction on the number of repetitions of α, simultaneously with the induction on the structure of hybrid games. This simultaneous induction is well-founded, because α^* only repeats α finitely often (least fixpoint). $s \in \varsigma_{\alpha^*}(X) = \bigcap\{Z \subseteq \mathcal{S} : X \cup \varsigma_\alpha(Z) \subseteq Z\}$ implies $s \in X$ or $s \in \varsigma_\alpha(\varsigma_{\alpha^*}(X))$. In the first case ($s \in X$), Angel already wins with the winning strategy $\sigma \overset{\text{def}}{=} \{(\mathsf{s})\}$, so we only need to consider the second case. By induction hypothesis (α is structurally simpler than α^*), this is equivalent to: $\sigma \overset{\text{def}}{=} \{(\mathsf{s})\}$ is a winning strategy for Angel for X from s or there is a winning strategy $\sigma_\alpha \subseteq \mathfrak{g}(\alpha)(s)$ for Angel, i.e., for all strategies $\tau \subseteq \mathfrak{g}(\alpha)(s)$ for Demon: Demon deadlocks or $\lceil \sigma_\alpha \oplus \tau \rceil_s \in \varsigma_{\alpha^*}(X)$. By induction hypothesis (from $\lceil \sigma_\alpha \oplus \tau \rceil_s$ Angel can win X with less repetitions than from s), $\lceil \sigma_\alpha \oplus \tau \rceil_s \in \varsigma_{\alpha^*}(X)$ is equivalent to the existence of a winning strategy σ_τ for Angel (which depends on the state $\lceil \sigma_\alpha \oplus \tau \rceil_s$ that the previous α game led to) with winning condition X from $\lceil \sigma_\alpha \oplus \tau \rceil_s$. This is equivalent to $\sigma \subseteq \mathfrak{g}(\alpha^*)(s)$ being a winning strategy for Angel for X from s, using

$$\sigma \overset{\text{def}}{=} \{(\mathsf{g})\} \cup \mathsf{g}\hat{}\sigma_\alpha \cup \bigcup_{\sigma_\alpha \oplus \tau} \mathsf{g}\hat{}(\sigma_\alpha \oplus \tau)\hat{}\sigma_\tau$$

The union is over all leaves $\sigma_\alpha \oplus \tau$ for which the game is not won by a player yet. Note that the above σ is a winning strategy for X, because, for all plays for which the game is decided during the first α, the strategy σ_α already wins the game. For the others, σ_τ wins the game from the respective state $\lceil \sigma_\alpha \oplus \tau \rceil_s$ that was reached by the actions $\sigma_\alpha \oplus \tau$ according to Demon's strategy τ for the first repetition of α. The converse direction uses the fact that every game play is finite, hence, all strategies choose g only finitely often on each path, which makes the repetition well-founded (least fixpoint).

7. $s \in \varsigma_{\alpha^d}(X) = \mathcal{S} \setminus \varsigma_\alpha(\mathcal{S} \setminus X) \iff s \notin \varsigma_\alpha(\mathcal{S} \setminus X)$. By induction hypothesis, this is equivalent to: there is no winning strategy $\sigma \subseteq \mathfrak{g}(\alpha)(s)$ for Angel winning $\mathcal{S} \setminus X$ from s. By Theorem 1, this is equivalent to: there is a winning strategy $\tau \subseteq \mathfrak{g}(\alpha)(s)$ for Demon winning X from s. Since the nodes where Angel acts swap with the nodes where Demon acts when moving from α to α^d, this is equivalent to: there is a winning strategy $\sigma \subseteq \mathfrak{g}(\alpha^d)(s)$ for Angel winning X from s using $\sigma \overset{\text{def}}{=} \{(\mathsf{d})\} \cup \mathsf{d}\hat{}\tau \cup \mathsf{d}\hat{}\tau\hat{}\mathsf{d}$.

□

D Soundness Proof

First, we prove that FP and ind are interderivable in the dGL calculus.

Proof of Lemma 2. Rule ind derives from FP: We first derive the following variant

$$(ind_R)\quad \frac{\psi \to [\alpha]\psi \quad \psi \to \phi}{\psi \to [\alpha^*]\phi}$$

From $\psi \to [\alpha]\psi$ and $\psi \to \phi$ propositionally derive $\psi \to \phi \wedge [\alpha]\psi$, from which contraposition and propositional logic yield $\neg\phi \vee \neg[\alpha]\psi \to \neg\psi$. By $[\alpha]\psi \equiv \neg\langle\alpha\rangle\neg\psi$, this is an abbreviation for $\neg\phi \vee \langle\alpha\rangle\neg\psi \to \neg\psi$. Now FP derives $\langle\alpha^*\rangle\neg\phi \to \neg\psi$, which, by duality, is $\neg[\alpha^*]\phi \to \neg\psi$, which gives $\psi \to [\alpha^*]\phi$ by contraposition. The classical []-induction rule ind follows by $\phi \overset{\text{def}}{\equiv} \psi$. From ind, the variant ind_R is derivable again by R on $\psi \to \phi$.

Rule FP derives from ind: From $\phi \vee \langle\alpha\rangle\psi \to \psi$, propositionally derive $\phi \to \psi$ and $\langle\alpha\rangle\psi \to \psi$. By R, the former gives $\langle\alpha^*\rangle\phi \to \langle\alpha^*\rangle\psi$. By contraposition, the latter derives $\neg\psi \to \neg\langle\alpha\rangle\psi$, which is $\neg\psi \to [\alpha]\neg\psi$ by duality. Now ind derives $\neg\psi \to [\alpha^*]\neg\psi$. By contraposition $\neg[\alpha^*]\neg\psi \to \psi$, which, by duality, is $\langle\alpha^*\rangle\psi \to \psi$. Thus, $\langle\alpha^*\rangle\phi \to \psi$ by the formula derived above. □

Now we prove soundness of the dGL proof calculus.

Proof of Theorem 3. Soundness of modus ponens (MP) is simple and not shown. In order to prove soundness of an implication axiom $\phi \to \psi$, we fix any set of states $\mathcal{S}$, and need to show $[\![\phi]\!] \subseteq [\![\psi]\!]$. To prove soundness of an equivalence axiom $\phi \leftrightarrow \psi$, we need to show $[\![\phi]\!] = [\![\psi]\!]$. To prove soundness of a rule

$$\frac{\phi}{\psi}$$

we consider any set of states $\mathcal{S}$ and assume that ϕ is valid in $\mathcal{S}$, i.e., $[\![\phi]\!] = \mathcal{S}$ and prove that ψ is valid in $\mathcal{S}$, i.e., $[\![\psi]\!] = \mathcal{S}$.

$\langle:=\rangle$ $[\![\langle x := \theta\rangle\phi(x)]\!] = \varsigma_{x:=\theta}([\![\phi(x)]\!]) = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in [\![\phi(x)]\!]\} = \{s \in \mathcal{S} : s \in [\![\phi(\theta)]\!]\} = [\![\phi(\theta)]\!]$, where the middle equation holds by the substitution lemma. We can use the classical substitution lemma if $\phi(\theta)$ is in first-order logic. Otherwise the proof of the substitution lemma for differential dynamic logic dL [Pla10, Lemma 2.2] immediately generalizes to dGL.

$\langle'\rangle$ $[\![\langle x' = \theta\rangle\phi]\!] = \varsigma_{x'=\theta}([\![\phi]\!]) = \{\varphi(0) \in \mathcal{S} : \frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ for all $\zeta \leq r$ for some $\varphi : [0, r] \to \mathcal{S}$ such that $\varphi(r) \in [\![\phi]\!]\}$. On the other hand, we have

$[\![\exists t{\geq}0\, \langle x := y(t)\rangle\phi]\!] = \{s \in \mathcal{S} : s_t^r \in [\![\langle x := y(t)\rangle\phi]\!]$ for some $r \geq 0\} = \{s \in \mathcal{S} : s_t^r \in \{u \in \mathcal{S} : u_x^{[\![y(t)]\!]_u} \in [\![\phi]\!]\}$ for some $r \geq 0\} = \{s \in \mathcal{S} : (s_t^r)_x^{[\![y(t)]\!]_{s_t^r}} \in [\![\phi]\!]$ for some $r \geq 0\}$.

The inclusion "$\supseteq$" between those two sides follows, because the function $\varphi(\zeta) := (s_t^\zeta)_x^{[\![y(t)]\!]_{s_t^\zeta}}$ solves the differential equation $x' = \theta$ by assumption. The inclusion "$\subseteq$" follows, because the solution of the smooth differential equation $x' = \theta$ is unique [Pla10, Lemma 2.1].

$\langle?\rangle$ $[\![\langle?\psi\rangle\phi]\!] = \varsigma_{?\psi}([\![\phi]\!]) = [\![\psi]\!] \cap [\![\phi]\!] = [\![\psi \wedge \phi]\!]$

$\langle\cup\rangle$ $[\![\langle\alpha \cup \beta\rangle\phi]\!] = \varsigma_{\alpha\cup\beta}([\![\phi]\!]) = \varsigma_\alpha([\![\phi]\!]) \cup \varsigma_\beta([\![\phi]\!]) = [\![\langle\alpha\rangle\phi]\!] \cup [\![\langle\beta\rangle\phi]\!] = [\![\langle\alpha\rangle\phi \vee \langle\beta\rangle\phi]\!]$

$\langle;\rangle$ $[\![\langle\alpha;\beta\rangle\phi]\!] = \varsigma_{\alpha;\beta}([\![\phi]\!]) = \varsigma_\alpha(\varsigma_\beta([\![\phi]\!])) = \varsigma_\alpha([\![\langle\beta\rangle\phi]\!]) = [\![\langle\alpha\rangle\langle\beta\rangle\phi]\!]$.

$\langle^d\rangle$ $[\![\langle\alpha^d\rangle\phi]\!] = \varsigma_{\alpha^d}([\![\phi]\!]) = \mathcal{S} \setminus \varsigma_\alpha(\mathcal{S} \setminus [\![\phi]\!]) = \mathcal{S} \setminus \varsigma_\alpha([\![\neg\phi]\!]) = \mathcal{S} \setminus [\![\langle\alpha\rangle\neg\phi]\!] = [\![\neg\langle\alpha\rangle\neg\phi]\!]$

$\langle^*\rangle$ Since $[\![\langle\alpha^*\rangle\phi]\!] = \varsigma_{\alpha^*}([\![\phi]\!]) = \mu Z.[\![\phi]\!] \cup \varsigma_\alpha(Z)$ is a fixpoint, we know that $[\![\langle\alpha^*\rangle\phi]\!] = [\![\phi]\!] \cup \varsigma_\alpha([\![\langle\alpha^*\rangle\phi]\!])$. Thus, $[\![\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi]\!] = [\![\phi]\!] \cup [\![\langle\alpha\rangle\langle\alpha^*\rangle\phi]\!] = [\![\phi]\!] \cup \varsigma_\alpha([\![\langle\alpha^*\rangle\phi]\!]) = [\![\langle\alpha^*\rangle\phi]\!]$. In particular, $[\![\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi]\!] \subseteq [\![\langle\alpha^*\rangle\phi]\!]$.

R Assume the premise $\phi \to \psi$ is valid in a state space $\mathcal{S}$, i.e., $[\![\phi]\!] \subseteq [\![\psi]\!]$. Then the conclusion $\langle\alpha\rangle\phi \to \langle\alpha\rangle\psi$ is valid in $\mathcal{S}$, i.e., $[\![\langle\alpha\rangle\phi]\!] = \varsigma_\alpha([\![\phi]\!]) \subseteq \varsigma_\alpha([\![\psi]\!]) = [\![\langle\alpha\rangle\psi]\!]$ by monotonicity of $\varsigma_\alpha(\cdot)$.

FP Assume the premise $\phi \vee \langle\alpha\rangle\psi \to \psi$ is valid in a state space $\mathcal{S}$, i.e., $[\![\phi \vee \langle\alpha\rangle\psi]\!] \subseteq [\![\psi]\!]$. Thus, $[\![\phi]\!] \cup \varsigma_\alpha([\![\psi]\!]) = [\![\phi]\!] \cup [\![\langle\alpha\rangle\psi]\!] = [\![\phi \vee \langle\alpha\rangle\psi]\!] \subseteq [\![\psi]\!]$. That is, ψ is a pre-fixpoint of $Z = [\![\phi]\!] \cup \varsigma_\alpha(Z)$. Now $[\![\langle\alpha^*\rangle\phi]\!] = \varsigma_{\alpha^*}([\![\phi]\!]) = \mu Z.[\![\phi]\!] \cup \varsigma_\alpha(Z)$ is the least fixpoint and even the least pre-fixpoint [Koz06, Appendix A], because of monotonicity. This implies $[\![\langle\alpha^*\rangle\phi]\!] \subseteq [\![\psi]\!]$, which implies that $\langle\alpha^*\rangle\phi \to \psi$ is valid in $\mathcal{S}$.

con By premise, we know for all values of v that $[\![\varphi(v) \wedge v > 0]\!] \subseteq [\![\langle\alpha\rangle\varphi(v-1)]\!] = \varsigma_\alpha([\![\varphi(v-1)]\!])$. To prove the conclusion, we show that for all values of v: $[\![\varphi(v)]\!] \subseteq [\![\langle\alpha^*\rangle\exists v{\leq}0\,\varphi(v)]\!] = \varsigma_{\alpha^*}([\![\exists v{\leq}0\,\varphi(v)]\!]) = \mu Z.[\![\exists v{\leq}0\,\varphi(v)]\!] \cup \varsigma_\alpha(Z)$. Since $[\![\exists v{\leq}0\,\varphi(v)]\!] \subseteq \varsigma_{\alpha^*}([\![\exists v{\leq}0\,\varphi(v)]\!])$, this holds trivially whenever $v \leq 0$. By induction on $r \in \mathbb{R}$, we assume $[\![\varphi(v)]\!] \subseteq \varsigma_{\alpha^*}([\![\exists v{\leq}0\,\varphi(v)]\!])$ for all $v \leq r$ and prove it for any $v > r$. It is enough to consider the case where $v < r+1$. Consider any $s \in [\![\varphi(v)]\!]$ (if no such s exists, there is nothing to show). Since $v > r \geq 0$, we know, by premise, that

$$s \in [\![\varphi(v) \wedge v > 0]\!] \subseteq \varsigma_\alpha([\![\varphi(v-1)]\!]) \overset{\text{mon}}{\subseteq} \varsigma_\alpha(\varsigma_{\alpha^*}([\![\exists v{\leq}0\,\varphi(v)]\!])) \overset{\mu}{\subseteq} \varsigma_{\alpha^*}([\![\exists v{\leq}0\,\varphi(v)]\!])$$

where the indicated inclusions are, respectively, by the induction hypothesis ($v - 1 \leq r$) together with monotonicity (mon) and the fact (marked μ) that $\varsigma_{\alpha^*}([\![\exists v{\leq}0\,\varphi(v)]\!])$ is a fixpoint. Thus, $s \in \varsigma_{\alpha^*}([\![\exists v{\leq}0\,\varphi(v)]\!]) = [\![\langle\alpha^*\rangle\exists v{\leq}0\,\varphi(v)]\!]$.

$\overleftarrow{\text{B}}$ We show that $[\![\exists x\,\langle\alpha\rangle\phi]\!] = \{s \in \mathcal{S} : s^r_x \in [\![\langle\alpha\rangle\phi]\!] = \varsigma_\alpha([\![\phi]\!])$ for some $r \in \mathbb{R}\}$ is contained in the following set, because of the assumption $x \notin \alpha$:

$[\![\langle\alpha\rangle\exists x\,\phi]\!] = \varsigma_\alpha([\![\exists x\,\phi]\!]) = \varsigma_\alpha(\{s \in \mathcal{S} : s^r_x \in [\![\phi]\!]$ for some $r \in \mathbb{R}\})$. Let $s \in \mathcal{S}$ with $s^r_x \in \varsigma_\alpha([\![\phi]\!])$ for some $r \in \mathbb{R}$. Since $x \notin \alpha$, $\varsigma_\alpha([\![\phi]\!])$ is independent of r, and the same sequence of game actions is applicable from s^r_x as from s. By $s^r_x \in \varsigma_\alpha([\![\phi]\!])$, there is a play of game α from s^r_x to some state of the form $t^r_x \in [\![\phi]\!]$. Note that x is unchanged during α. Without loss of generality, we can choose t to be a state with $t(x) = r$. Since $x \notin \alpha$, the exact same play of game α leads from s to t, just with the value $s(x)$ for x. This proves the inclusion "$\subseteq$" of the above sets, because $t \in [\![\exists x\,\phi]\!]$. Note that the inclusion "$\supseteq$" does not hold, because, even if $x \notin \alpha$, the winning states in the second set depend on the value of x, so the strategy may depend on that value. □